On the quantum hardness of solving isomorphism problems as nonabelian hidden shift problems

Andrew M. Childs[*] and Paweł Wocjan[†]

Institute for Quantum Information
California Institute of Technology
Pasadena, CA 91125, USA

Abstract

We consider an approach to deciding isomorphism of rigid n-vertex graphs (and related isomorphism problems) by solving a nonabelian hidden shift problem on a quantum computer using the standard method. Such an approach is arguably more natural than viewing the problem as a hidden subgroup problem. We prove that the hidden shift approach to rigid graph isomorphism is hard in two senses. First, we prove that $\Omega(n)$ copies of the hidden shift states are necessary to solve the problem (whereas $O(n \log n)$ copies are sufficient). Second, we prove that if one is restricted to single-register measurements, an exponential number of hidden shift states are required.

1 Introduction

One of the major challenges of quantum computing is to determine whether there exists an efficient quantum algorithm to decide if two graphs are isomorphic. It is well known that the graph isomorphism problem can be reduced to a hidden subgroup problem over the symmetric group [5, 7, 10, 20]. This approach seems to be promising since hidden subgroup problems over many groups, including arbitrary abelian groups [7, 16, 23, 32, 33] and some nonabelian ones [4, 13–15, 21, 25] can be solved efficiently on a quantum computer. An efficient quantum algorithm for graph isomorphism would be interesting since no efficient classical algorithm for the problem is known; the best known classical algorithm for deciding isomorphism of n-vertex graphs runs in time $O(n^{\sqrt{cn/\log n}})$ for some constant c [2].

Unfortunately, the only results so far on the quantum complexity of the graph isomorphism problem consist of evidence that the problem might be hard (with the notable exception of the result that the query complexity of the associated hidden subgroup problem is polynomial [10]). The graph isomorphism problem can be reduced to a hidden subgroup problem in S_{2n} where the hidden subgroups are generated by full support involutions. Hallgren, Russell, and Ta-Shma showed that *weak Fourier sampling*, in which one performs a nonabelian Fourier transform but then only measures the name

---

[*]amchilds@caltech.edu

[†]wocjan@cs.caltech.edu

of an irreducible representation, is insufficient to solve the problem [18]. Kempe and Shalev generalized their result to show that finding other subgroups of the symmetric group is also hard [22]. Finally, Moore, Russell, and Schulman have obtained results about the need to use multi-register measurements on the hidden subgroup states obtained by Fourier sampling. In particular, if one is restricted to single-register measurements (in the standard approach known as *strong Fourier sampling*), an exponential number of hidden subgroup states is required [27]. Similarly, if one is restricted to two-register measurements, then a superpolynomial (though possibly subexponential) number of hidden subgroup states is required [26]. Strictly speaking, these results do not show that the hidden subgroup problem directly relevant to graph isomorphism is hard, since the possible subgroups resulting from the graph isomorphism reduction are not generated by arbitrary full support involutions, but by involutions having further properties (as we will discuss further in Section 3, in connection with the hidden subgroup problem over $S_n \wr \mathbb{Z}_2$). However, concurrently with the present work, Moore, Russell, and Schulman have improved their result for single-register measurements to cover the special case directly relevant to graph isomorphism [28].

In this paper, we study an alternative approach to solving graph isomorphism on a quantum computer, by viewing it as an instance of a nonabelian hidden *shift* problem. This approach is arguably more natural than viewing the problem as a hidden subgroup problem: every possible hidden shift corresponds to a possible isomorphism (whereas there are many subgroups of either S_{2n} or $S_n \wr \mathbb{Z}_2$ that do not correspond to isomorphisms); and furthermore, viewed as black box problems, the hidden shift problem can be reduced to the hidden subgroup problem. The hidden shift problem can be tackled on a quantum computer using a standard method that closely parallels the standard approach to the hidden subgroup problem. We present two hardness results for this standard approach to the hidden shift problem over S_n.

First, we prove that $\Omega(n)$ copies of the hidden shift state are necessary to solve the problem (whereas $O(n \log n)$ copies are sufficient). The idea behind this bound is the simple observation that the hidden shift problem for the largest abelian subgroup of S_n is at least as hard as for the whole group S_n. In the case where the group G is abelian, the hidden shift problem for G is equivalent to the hidden subgroup problem over the generalized dihedral group $G \rtimes \mathbb{Z}_2$, and it is straightforward to obtain a reasonably tight bound for this case using a connection to the subset sum problem over G. Since S_n contains large abelian subgroups, the resulting bound for the nonabelian hidden shift problem is not too bad.

Second, we present a simple proof that single-register measurements are not sufficient to solve the hidden shift problem over S_n. In fact, this result holds for any group that has many irreducible representations of sufficiently high degree. In particular, the only property of S_n used in the proof is the fact that under the Plancherel distribution, an irreducible representation of S_n with degree larger than $n^{\Theta(n)}$ occurs with probability at least $1 - n^{-\Omega(n)}$.

The remainder of the paper is organized as follows. In Section 2 we define the nonabelian hidden shift problem and discuss the standard approach to solving it. In Section 3 we discuss how isomorphism problems (including, but not limited to, graph isomorphism) can be cast as hidden shift problems. In Section 4 we prove the linear lower bound on the required number of copies of hidden shift states. In Section 5 we examine the structure of the hidden shift states for arbitrary groups and obtain some results needed for Section 6, where we show that single register measurements are

insufficient. Finally, in the Appendix, we present some additional results on the rank of the hidden shift states.

2 Nonabelian hidden shift problem

The (nonabelian) hidden shift problem is the following. We are given black-box access to two functions $f_0 : G \to S$ and $f_1 : G \to S$ where G is a (nonabelian) group and S is a finite set. The functions are promised to satisfy two conditions:

1. Both f_0 and f_1 are injective.
2. Either there exists a fixed hidden shift $s \in G$ such that $f_0(g) = f_1(gs)$ for all $g \in G$, or the images of f_0 and f_1 are disjoint (in which case we say there is no hidden shift).

The goal is to determine whether there is a hidden shift s or not.

The case where G is an abelian group has received considerable attention [3, 9, 11, 13, 24, 30, 31]. Since inversion is an automorphsim of any abelian group, the hidden shift problem in G is a hidden subgroup problem in the generalized dihedral group $G \rtimes \mathbb{Z}_2$ where $\mathbb{Z}_2$ acts by inversion. In particular, the case where G is cyclic is the well-known dihedral hidden subgroup problem. However, the case where G is nonabelian, in which case the hidden shift problem is not a hidden subgroup problem, seems not to have been studied extensively.

In this paper, we focus on a particular natural approach to solving the hidden shift problem on a quantum computer, paralleling the standard quantum approach to the hidden subgroup problem. First prepare a uniform superposition over $i \in \mathbb{Z}_2$ and $g \in G$, and then compute the value of $f_i(g)$, giving the state

$$\frac{1}{\sqrt{2|G|}} \sum_{g \in G} \left(|0, g, f_0(g)\rangle + |1, g, f_1(g)\rangle \right). \tag{1}$$

Then measure the third register. If there is a hidden shift s, then we are left with the state

$$|\phi_{s,g}\rangle := \frac{1}{\sqrt{2}}(|0, g\rangle + |1, gs\rangle) \tag{2}$$

for some uniformly random (unknown) $g \in G$. On the other hand, if there is no hidden shift, we obtain the state $|i, g\rangle$ for some uniformly random (unknown) $i \in \mathbb{Z}_2$ and $g \in G$. Thus the density matrix obtained by applying the procedure is either

$$\gamma_1(s) := \frac{1}{|G|} \sum_{g \in G} |\phi_{s,g}\rangle\langle\phi_{s,g}| \tag{3}$$

if there is a hidden shift $s \in G$, or the maximally mixed state

$$\gamma_2 := \frac{1}{2|G|} I_{2|G|} \tag{4}$$

if there is no hidden shift. Using the state thus obtained, we would like to decide whether there is a hidden shift or not.

In general, we can apply the above procedure k times to obtain k copies of the hidden shift state (or the maximally mixed state if there is no hidden shift). Clearly, these states become more distinguishable as k is increased. Suppose that in the case where there is a hidden shift s, it is equally likely to correspond to any element of G. Then the problem is to distinguish the two density operators

$$\gamma_1^{(k)} := \frac{1}{|G|} \sum_{s \in G} \gamma_1^{(k)}(s) \tag{5}$$

$$\gamma_2^{(k)} := \frac{1}{(2|G|)^k} I, \tag{6}$$

where $\gamma_1^{(k)}(s) := \gamma_1(s)^{\otimes k}$.

A natural generalization of the nonabelian hidden shift problem involves the case of M injective functions, f_j for $j \in \{0, 1, \dots, M-1\}$, satisfying $f_j(g) = f_{j+1}(gs)$ for a fixed $s \in G$ for all $j \in \{0, 1, \dots, M-2\}$. This problem becomes easier as M is increased, and is interesting in the case where G is cyclic, since it has an efficient quantum algorithm provided M is sufficiently large [8]. We will not consider the generalized nonabelian hidden shift problem further in this paper, although it is an interesting question whether this problem has an efficient quantum algorithm even for M sufficiently large.

3   Isomorphism problems

The nonabelian hidden shift problem for the symmetric group is especially interesting since an efficient quantum algorithm for this problem would yield an efficient algorithm for graph isomorphism (and more generally, for other related isomorphism problems). The usual quantum approach to graph isomorphism relies on a reduction to the hidden subgroup problem for the symmetric group, but the hidden shift problem for $G = S_n$ presents an alternative approach that seems to be at least as natural, and is arguably more so.

We now describe a generalized isomorphism problem that reduces to the hidden shift problem. For each $n \in \mathbb{N}$, let $\mathcal{C}_n$ be a set of objects of size n. For example, $\mathcal{C}_n$ could be the set of graphs on n vertices. We assume that the objects can be uniquely represented using $\mathrm{poly}(n)$ bits.

Let G_n be a family of (finite) groups such that each G_n acts on $\mathcal{C}_n$. For $g \in G_n$ and $C \in \mathcal{C}_n$, let $g(C)$ denote the element of $\mathcal{C}_n$ given by the action of g on C. We call two objects $A, B \in \mathcal{C}_n$ *isomorphic* if there is some $g \in G_n$ such that $g(A) = B$. We call an object $C \in \mathcal{C}_n$ *rigid* if it has no automorphisms, i.e., if there is no $g \in G_n - \{1\}$ such that $g(C) = C$.

The $\mathcal{C}$-*isomorphism problem* is the following. Given two rigid objects $C_0, C_1 \in \mathcal{C}_n$, determine whether they are isomorphic or nonisomorphic. It is straightforward to reduce this isomorphism problem to a corresponding hidden shift problem: simply let $f_i(g) := g(C_i)$. (The assumption of rigidity is required to ensure that f_0, f_1 are injective.)

Graph isomorphism is the special case of the $\mathcal{C}$-isomorphism problem for $G_n = S_n$ where $\mathcal{C}_n$ is the set of graphs on n vertices, and the action of G_n is to permute the vertices. Thus, a solution to the generalized hidden shift problem for $G_n = S_n$ would give an efficient algorithm for testing isomorphism of rigid graphs. But such an algorithm could also be applied to other isomorphism problems. For example, if we let

$\mathcal{C}_n$ be the set of all binary linear codes of length n, where $G_n = S_n$ acts to permute the bits of the code words, then we obtain the *code equivalence problem* discussed in [10], which is at least as hard as graph isomorphism [29].

As mentioned in the introduction, the usual approach to solving graph isomorphism on a quantum computer is based not on the hidden shift problem, but on the hidden subgroup problem. Graph isomorphism can be cast as a hidden subgroup problem over S_{2n} where the hidden subgroups are generated by full support involutions. A more careful inspection of the hidden subgroups that occur in this reduction shows that it is sufficient work with a subgroup of S_{2n}: as proposed by Ettinger and Høyer, one can cast graph isomorphism as a hidden subgroup problem over the wreath product $S_n \wr \mathbb{Z}_2 < S_{2n}$ where the hidden subgroups are generated by so-called involutive swaps [10].

How are the hidden subgroup and hidden shift approaches to graph isomorphism related? In general, one can show that the hidden shift problem in a group G can be reduced to the hidden subgroup problem in $G \wr \mathbb{Z}_2$. In particular, the hidden shift problem in S_n reduces to the hidden subgroup problem in $S_n \wr \mathbb{Z}_2$ (and in fact, using the results of [36], one can also show that it reduces to the hidden subgroup problem in S_{2n}). Since the hidden shift problem is no harder than the corresponding hidden subgroup problem, this suggests that the hidden shift problem might present a more natural quantum approach to graph isomorphism. However, we emphasize that our hardness results about measurements of hidden shift states do not imply corresponding results about hidden subgroup states, since the reduction does not necessarily still hold when we assume the use of the standard method to produce particular quantum states.

4   Lower bound on the number of copies

In this section, we show that $\Omega(n)$ copies of the hidden shift states are needed to successfully determine whether there is a hidden shift. We do this by showing that the optimal POVM is unlikely to produce the correct answer unless $k = \Omega(n)$.

Consider the general problem of distinguishing a pair of (possibly mixed, a priori equiprobable) quantum states. The optimal measurement for this problem (in the sense that it maximizes the probability of successfully identifying the state) was discovered by Helstrom [19], and is as follows. Suppose we wish to distinguish the quantum states ρ_1, ρ_2. Then let E_1 be the projector onto the eigenvectors of $\rho_1 - \rho_2$ corresponding to positive eigenvalues, and let E_2 be the projector onto the eigenvectors of $\rho_1 - \rho_2$ corresponding to negative eigenvalues. (Eigenvectors in the nullspace of $\rho_1 - \rho_2$ can be associated to either E_1 or E_2 without affecting the success probability.)

In principle, Helstrom's result tells us the optimal measurement to distinguish $\gamma_1^{(k)}$ and $\gamma_2^{(k)}$. Unfortunately, since we do not have a good understanding of the spectrum of $\gamma_1^{(k)}$ for nonabelian groups, we do not know how to estimate the success probability of the Helstrom measurement in such cases. However, we can obtain a good estimate of the success probability for abelian groups, and we can obtain a bound for arbitrary groups since a bound for a subgroup implies a bound for the full group. Specifically, we have

Lemma 1. *The number of copies needed to solve the hidden shift problem in the group G (with a probability of success bounded above $1/2$ by a constant) is at least as great as the number of copies needed to solve the hidden shift problem in any subgroup $H \le G$.*

Proof. Clearly, if the possible hidden shifts are restricted to be from a subgroup $H \leq G$, the problem is at least as hard as when the hidden shift may be arbitrary. For a uniformly random hidden shift $s \in H$, the density matrix when there is a hidden shift is

$$\frac{1}{|H|} \sum_{h \in H} |\phi_{s,h}\rangle\langle\phi_{s,h}|, \tag{7}$$

which can be written as the tensor product of the unrestricted hidden shift state in H and a maximally mixed state of dimension $|G|/|H|$. Since the maximally mixed state provides no information about the hidden shift, the restricted problem in G is equivalent to the hidden shift problem for H. $\qquad\square$

Now we give a general lower bound on the number of copies needed to solve an arbitrary abelian hidden shift problem. In the abelian case, we can give fairly tight bounds using the close connection between the hidden shift problem and the subset sum problem [3]. Specifically, after performing a Fourier transform on the group register, we can write the abelian hidden shift states as

$$\tilde{\gamma}_1^{(k)}(s) = \frac{1}{(2|G|)^k} \sum_{x \in G^k} \sum_{w,v \in G} \chi_w(s)\bar{\chi}_v(s)\sqrt{\eta_w^x \eta_v^x}|S_w^x, x\rangle\langle S_v^x, x| \tag{8}$$

where

$$S_w^x := \{b \in \mathbb{Z}_2^k : b \cdot x = w\} \tag{9}$$

is the set of solutions of the subset sum problem over G, $\eta_w^x := |S_w^x|$ is the number of such solutions, and

$$|S_w^x\rangle := \frac{1}{\sqrt{\eta_w^x}} \sum_{b \in S_w^x} |b\rangle \tag{10}$$

is the normalized uniform superposition over those solutions (where we define $|S_w^x\rangle := 0$ in the event that $\eta_w^x = 0$). Thus, with a uniformly random hidden shift, we have the state

$$\tilde{\gamma}_1^{(k)} = \frac{1}{(2|G|)^k} \sum_{x \in G^k} \sum_{w \in G} \eta_w^x |x, S_w^x\rangle\langle x, S_w^x|. \tag{11}$$

In the standard approach to the abelian hidden shift problem, our goal is to distinguish this state from the maximally mixed state. An optimal measurement for doing so is the measurement that projects onto the support of $\tilde{\gamma}_1^{(k)}$. Since the eigenvalues of $\tilde{\gamma}_1^{(k)}$ are integer multiples of $1/(2|G|)^k$, the operator $\tilde{\gamma}_1^{(k)} - \tilde{\gamma}_2^{(k)}$ is nonnegative precisely on the support of $\tilde{\gamma}_1^{(k)}$. Therefore, the projection onto that support is a Helstrom measurement, and hence is optimal.

Having identified an optimal measurement, we can now show

Lemma 2. *For any abelian group G, $k = \Omega(\log|G|)$ copies of the hidden shift states are needed to decide whether there is a hidden shift (with a probability of success bounded above $1/2$ by a constant).*

Proof. The success probability of the optimal measurement (in which E_1 projects onto the support of $\tilde{\gamma}_1^{(k)}$ and E_2 projects onto its complement) is

$$\Pr(\text{success}) := \frac{1}{2}\left(\operatorname{tr} E_1\tilde{\gamma}_1^{(k)} + \operatorname{tr} E_2\tilde{\gamma}_2^{(k)}\right) \tag{12}$$

$$= 1 - \frac{\operatorname{rank}\tilde{\gamma}_1^{(k)}}{2(2|G|)^k}. \tag{13}$$

Now

$$\operatorname{rank}\tilde{\gamma}_1^{(k)} = \sum_{x,w}\delta[\eta_w^x > 0] \tag{14}$$

$$= |G|^{k+1} - \sum_{x,w}\delta[\eta_w^x = 0]. \tag{15}$$

(For the case $G = \mathbb{Z}_N$, the rank is given by the integer sequence [34, A098966]. For a discussion of the rank in the general (not necessarily abelian) case, see the Appendix.) To evaluate this expression, we need to understand the typical behavior of η_w^x. In particular, it is helpful to know the first and second moments of η_w^x for uniformly random $x \in G^k$, $w \in G$. For an arbitrary group G, the first moment is

$$\mu := \mathop{\mathbf{E}}_{x,w}\eta_w^x = \frac{2^k}{|G|}. \tag{16}$$

For the second moment, we have

$$\mathop{\mathbf{E}}_{x,w}(\eta_w^x)^2 := \frac{1}{|G|^{k+1}}\sum_{x,w}(\eta_w^x)^2 \tag{17}$$

$$= \frac{1}{|G|^{k+1}}\sum_{x,w}\left(\sum_b\delta_{b\cdot x,w}\right)^2 \tag{18}$$

$$= \mu + \frac{1}{|G|^{k+1}}\sum_{x,w}\sum_{b\neq c}\delta_{b\cdot x,c\cdot x}\,\delta_{b\cdot x,w} \tag{19}$$

$$= \mu + \frac{1}{|G|^{k+1}}\sum_{x}\sum_{b\neq c}\delta_{b\cdot x,c\cdot x} \tag{20}$$

$$= \mu + \frac{2^k(2^k-1)}{|G|^2}. \tag{21}$$

Here in the final step we used the fact that for fixed $b \neq c$ (with $b_k \neq c_k$ without loss of generality), and for fixed $x_1,\dots,x_{k-1} \in G$, there is exactly one $x_k \in G$ such that $b \cdot x = c \cdot x$. In terms of the variance $\sigma^2 := \mathbf{E}_{x,w}(\eta_w^x)^2 - \mu^2$ we have the inequality $\Pr(\eta_w^x = 0) \leq \sigma^2/(\mu^2 + \sigma^2)$ [1], giving

$$\operatorname{rank}\gamma_1^{(k)} \geq |G|^{k+1}\frac{\mu^2}{\mu^2+\sigma^2} \tag{22}$$

$$= |G|^{k+1}\left(\mu + 1 - \frac{1}{|G|}\right)^{-1} \tag{23}$$

$$\geq (2|G|)^k - |G|^{k+1}. \tag{24}$$

7

Thus, we find

$$\Pr(\text{success}) \le \frac{1}{2}\left(1 + \frac{|G|}{2^k}\right). \tag{25}$$

For the success probability to be bounded above $1/2$ by a constant, we need $k = \Omega(\log|G|)$ as claimed. $\qquad\square$

Putting these lemmas together, we have

Theorem 3. *To solve the hidden shift problem in S_n, $\Omega(n)$ copies of the hidden shift states are necessary.*

Proof. The largest abelian subgroup of S_n has size $3^{\Theta(n)}$ [6] (see also [34, A000792]). Combining Lemmas 1 and 2 gives the result. $\qquad\square$

This result is not too far from the best possible, since $O(\log|G|)$ copies are sufficient to solve the hidden shift problem for any group G. This follows easily from (13) and the fact that $\operatorname{rank}\tilde{\gamma}_1^{(k)} \le |G|^{k+1}$, and is analogous to the well-known result that $O(\log|G|)$ copies of hidden subgroup states are sufficient to solve the hidden subgroup problem [12]. However, there is a logarithmic gap between these lower and upper bounds. We suspect that the lower bound could be improved, since it only uses information about abelian subgroups, but without a better understanding of the structure of the hidden shift states for large k, it seems difficult to establish a bound.

Note that an analogous bound of $\Omega(n)$ has recently been independently established for the hidden subgroup problem over S_{2n} where the hidden subgroup may be an arbitrary full-support involution, and a bound of $\Omega(n\log n)$ has been established for the hidden subgroup problem over $S_n \wr \mathbb{Z}_2$ [17].

It is worth noting that while the projection onto the support of $\gamma_1^{(k)}$ is an optimal measurement in the abelian case, it is not an optimal measurement in general. For example, for $G = S_4$, $\gamma_1^{(3)}$ has eigenvalues between 0 and $1/(2|G|)^3$, so that the projection onto the support is not a Helstrom measurement.

5 Structure of hidden shift states

To show that single-register measurements are not sufficient to solve the hidden shift problem, we need to understand the structure of the states $\gamma_1^{(k)}(s)$, $\gamma_1^{(k)}$, and $\gamma_2^{(k)}$. Here we determine their block structure and use it to compute the spectrum of $\gamma_1^{(k)}$ for $k = 1$ and 2.

Observe that $\gamma_1(s)$ has the following form:

$$\gamma_1(s) = \frac{1}{2|G|}\sum_{g\in G} |0,g\rangle\langle0,g| + |1,gs\rangle\langle1,gs| + |0,g\rangle\langle1,gs| + |1,gs\rangle\langle0,g| \tag{26}$$

$$= \frac{1}{2|G|}\begin{pmatrix} I & R(s) \\ R(s^{-1}) & I \end{pmatrix}, \tag{27}$$

where R is the *right regular representation* of G, defined by

$$R(s)|g\rangle = |gs^{-1}\rangle \tag{28}$$

8

for all $s, g \in G$. Recall that the regular representation contains all irreducible representations of G with multiplicities given by their dimensions. More precisely, we have

$$F\,R(s)\,F^\dagger = \bigoplus_{\rho \in \hat{G}} I_{d_\rho} \otimes \rho(s) \tag{29}$$

for all $s \in G$, where F is the Fourier transform over G and $\hat{G}$ is a complete set of irreducible representations of G. In other words, the Fourier transform decomposes the regular representation into its irreducible constituents.

Using the Fourier transform, the states $\gamma_1^{(k)}(s)$, $\gamma_1^{(k)}$, and $\gamma_2^{(k)}$ can be simultaneously block diagonalized for any $k \in \mathbb{N}$. The blocks are enumerated by k-tuples of irreducible representations. In particular, in the Fourier basis we have

$$\tilde{\gamma}_1^{(k)}(s) = \frac{1}{(2|G|)^k} \bigoplus_{(\rho_1,\ldots,\rho_k) \in \hat{G}^k} I_{d_{\rho_1} \cdots d_{\rho_k}} \otimes B^{\rho_1,\ldots,\rho_k}(s) \tag{30}$$

$$\tilde{\gamma}_1^{(k)} = \frac{1}{(2|G|)^k} \bigoplus_{(\rho_1,\ldots,\rho_k) \in \hat{G}^k} I_{d_{\rho_1} \cdots d_{\rho_k}} \otimes B^{\rho_1,\ldots,\rho_k} \tag{31}$$

$$\tilde{\gamma}_2^{(k)} = \frac{1}{(2|G|)^k} \bigoplus_{(\rho_1,\ldots,\rho_k) \in \hat{G}^k} I_{d_{\rho_1} \cdots d_{\rho_k}} \otimes I_{2d_{\rho_1} \cdots 2d_{\rho_k}} \tag{32}$$

where

$$B^{\rho_1,\ldots,\rho_k}(s) := \bigotimes_{j=1}^{k} \begin{pmatrix} I_{d_{\rho_j}} & \rho_j(s) \\ \rho_j(s^{-1}) & I_{d_{\rho_j}} \end{pmatrix} \tag{33}$$

$$B^{\rho_1,\ldots,\rho_k} := \frac{1}{|G|} \sum_{s \in G} B^{\rho_1,\ldots,\rho_k}(s)\,. \tag{34}$$

Here the factor $d_{\rho_1} \cdots d_{\rho_k}$ accounts for the multiplicity of $(\rho_1,\ldots,\rho_k)$ in k copies of the regular representation of G.

It is straightforward to check that the blocks $B^{\rho_1,\ldots,\rho_k}(s)$ and $B^{\rho_1,\ldots,\rho_k}$ can be expressed as

$$B^{\rho_1,\ldots,\rho_k}(s) = \sum_{x,y \in \{0,1\}^k} |x\rangle\langle y| \otimes A^{\rho_1,\ldots,\rho_k}_{y_1-x_1,\ldots,y_k-x_k}(s) \tag{35}$$

$$B^{\rho_1,\ldots,\rho_k} = \sum_{x,y \in \{0,1\}^k} |x\rangle\langle y| \otimes A^{\rho_1,\ldots,\rho_k}_{y_1-x_1,\ldots,y_k-x_k}\,, \tag{36}$$

where

$$A^{\rho_1,\ldots,\rho_k}_{z_1,\ldots,z_k}(s) := \rho_1(s^{z_1}) \otimes \rho_2(s^{z_2}) \otimes \cdots \otimes \rho_k(s^{z_k}) \tag{37}$$

$$A^{\rho_1,\ldots,\rho_k}_{z_1,\ldots,z_k} := \frac{1}{|G|} \sum_{s \in G} A^{\rho_1,\ldots,\rho_k}_{z_1,\ldots,z_k}(s) \tag{38}$$

for all $z \in \{-1,0,1\}^k$. Clearly, the matrices $A^{\rho_1,\ldots,\rho_k}_{z_1,\ldots,z_k}$ are hermitian, that is, $A^{\rho_1,\ldots,\rho_k}_{z_1,\ldots,z_k} = A^{\rho_1,\ldots,\rho_k}_{-z_1,\ldots,-z_k}$.

To understand the form of these matrices, we must carry out the sum in (38) for various choices of the irreducible representations $\rho_1, \ldots, \rho_k \in \hat{G}$ and the indices $z_1, \ldots, z_k \in \{-1, 0, 1\}$. If all z_j have the same sign, then such a calculation is straightforward, using the following well-known result:

Lemma 4. *Let π be any representation of the group G. Then the matrix*

$$A := \frac{1}{|G|} \sum_{g \in G} \pi(g) \tag{39}$$

is a projection operator whose rank is the number of times the trivial representation appears in π.

Proof. Decompose the representation π into irreducible representations. Let σ be any irreducible representation occurring in π. The sum $B = \frac{1}{|G|} \sum_{g \in G} \sigma(g)$ is a multiple of the identity matrix because B commutes with all $\sigma(h)$ for $h \in G$. The trace of B is the inner product of the trivial character and the character of σ. Therefore, $B = 1$ if σ is the trivial representation and B is the zero matrix if σ is not the trivial representation. $\quad\square$

In general, we will have z_j's of both signs. In this case we may say that A includes both representations and antirepresentations of G, since $g \mapsto \rho(g^{-1})$ is a group antihomomorphism. Fortunately, this case can be dealt with using the following:

Lemma 5. *Let ρ and σ be two irreducible representations of G. Then the entries of the matrix*

$$A := \frac{1}{|G|} \sum_{g \in G} \rho(g) \otimes \sigma(g^{-1}) \tag{40}$$

are given by

$$A_{i,j;k,l} = \delta_{\rho,\sigma} \frac{1}{d_\rho} \delta_{i,l} \, \delta_{j,k} \tag{41}$$

where i, j are the row and column indices of the first tensor component and k, l are the row and column indices of the second tensor component.

Proof. The entries are given by

$$A_{i,j;k,l} = \frac{1}{|G|} \sum_{g \in G} \rho_{ij}(g) \otimes \overline{\sigma_{lk}(g)} \, ; \tag{42}$$

then (41) follows directly from the Schur orthogonality relations. $\quad\square$

Now we are ready to investigate the blocks B^ρ for $k = 1$ and the blocks $B^{\rho,\rho}$ for $k = 2$.

Lemma 6 (Spectrum for $k = 1$). *The block $B^{\hat{1}}$ has eigenvalues 2 and 0. For $\rho \neq \hat{1}$, $B^\rho = I_{2d_\rho}$.*

Proof. Since $\rho_{\hat{1}}(s) = 1$ for all s,

$$B^{\hat{1}} = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, \tag{43}$$

which has eigenvalues 2 and 0. For $\rho \neq \hat{1}$, $\sum_{s \in G} \rho(s) = 0$ by the orthogonality of ρ and $\hat{1}$, so that $B^\rho = I_{2d_\rho}$ as claimed. $\quad\square$

10

Lemma 7 (Spectrum for $k = 2$). *For any $\rho \in \hat{G} - \{\hat{1}\}$, either $B^{\rho,\rho}$ has the spectrum 1 (multiplicity $2d_\rho^2$) and $1 \pm 1/d_\rho$ (multiplicity d_ρ^2 each); or the spectrum 2 (multiplicity 1), 0 (multiplicity 1), 1 (multiplicity $2d_\rho^2 - 2$), and $1 \pm 1/d_\rho$ (multiplicity d_ρ^2 each).*

Proof. For simplicity, we omit the label ρ, ρ. The block of interest has the form

$$
B = \begin{pmatrix} I & A_{0,1} & A_{1,0} & A_{1,1} \\ A_{0,1} & I & A_{1,-1} & A_{1,0} \\ A_{1,0} & A_{1,-1} & I & A_{0,1} \\ A_{1,1} & A_{1,0} & A_{0,1} & I \end{pmatrix}. \tag{44}
$$

Recall that the blocks of B are enumerated by $x, y \in \{0,1\}^2$. The matrix at position (x, y) is given by A_{y-x}, where the A matrices are defined by

$$
A_z := \frac{1}{|G|} \sum_{g \in G} \rho(g^{z_1}) \otimes \rho(g^{z_2}) \tag{45}
$$

for $z \in \{-1, 0, 1\}^2$. We have simplified (44) to minimize the number of -1's using the fact that A_z is hermitian, so $A_z = A_{-z}$.

Since $\rho \neq \hat{1}$ by assumption, $A_{0,1} = A_{1,0} = 0$ by the calculation in Lemma 6. Thus

$$
B = \begin{pmatrix} I & 0 & 0 & A_{1,1} \\ 0 & I & A_{1,-1} & 0 \\ 0 & A_{1,-1} & I & 0 \\ A_{1,1} & 0 & 0 & I \end{pmatrix} \cong \begin{pmatrix} I & A_{1,1} \\ A_{1,1} & I \end{pmatrix} \oplus \begin{pmatrix} I & A_{1,-1} \\ A_{1,-1} & I \end{pmatrix}. \tag{46}
$$

Hence it remains to understand the operators $A_{1,1}$ and $A_{1,-1}$.

Since ρ (and hence also $\bar{\rho}$) is irreducible, the trivial representation appears at most once in $\rho \otimes \rho$, so by Lemma 4, $A_{1,1}$ is either zero or a projector of rank one. Hence the matrix

$$
\begin{pmatrix} I & A_{1,1} \\ A_{1,1} & I \end{pmatrix} \tag{47}
$$

is either the identity, or has the eigenvalues 2 and 0 with multiplicity 1, and 1 with multiplicity $2d_\rho^2 - 1$. By Lemma 5, $A_{1,-1}$ has eigenvalues $\pm 1/d_\rho$, so that

$$
\begin{pmatrix} I & A_{1,-1} \\ A_{1,-1} & I \end{pmatrix} \tag{48}
$$

has the eigenvalues $1 \pm 1/d_\rho$ each with multiplicity d_ρ^2. $\qquad \square$

6 Single-register measurements do not suffice

In this section we show that single-register measurements do not suffice to efficiently solve the hidden shift problem for $G = S_n$.

Let us first explain in more detail what is meant by an algorithm restricted to single-register measurements. A POVM $\mathcal{E}$ with a set of possible outcomes J is a collection of positive operators $\mathcal{E} = \{E_j : j \in J\}$ satisfying the completeness condition

$$
\sum_j E_j = I. \tag{49}
$$

An efficient algorithm consists of a polynomial number of POVMs $\mathcal{E}_1, \ldots, \mathcal{E}_t$, each acting on a single copy of the hidden shift state. After obtaining the measurement outcomes $j_1, \ldots, j_t$, a final computation is performed to decide whether there is a hidden shift or not. Note that the individual outcomes j_i need not directly correspond to one situation or the other. Also, let us stress out that the POVMs $\mathcal{E}_1, \ldots, \mathcal{E}_t$ may be chosen adaptively, that is, $\mathcal{E}_r$ may depend on all previous outcomes $j_1, \ldots, j_{r-1}$ for $2 \le r \le t$.

To simplify the analysis, we can refine any POVM $\mathcal{E}$ so that each $E_j = a_j |\psi_j\rangle\langle\psi_j|$ where each $|\psi_j\rangle$ is a unit vector and $a_j > 0$ without loss of generality. This is because any positive operator can be written as a weighted sum of projection operators, where the weights correspond to the eigenvalues and the projection operators to the eigenspaces. The result of this measurement on the state γ is a random variable, where we obtain $j \in J$ with probability

$$p(j) = a_j \langle\psi_j|\gamma|\psi_j\rangle. \tag{50}$$

In our case, the POVM can be further simplified because the states $\gamma_1^{(k)}(s)$, $\gamma_1^{(k)}$, and $\gamma_2^{(k)}$ can be simultaneously block-diagonalized as described in Section 5. The blocks are labeled by irreducible representations of G. Therefore, as in the hidden subgroup problem, we may assume without loss of generality that we first perform a Fourier transform on the group register and then measure the representation name (so-called *weak Fourier sampling*). Next, we perform a measurement within the subspace corresponding to the observed representation.

From the block decomposition of the states described in Section 5, it is clear that the various irreducible representations of G occur independently according to the Plancherel distribution, i.e.,

$$\Pr(\rho) = \frac{d_\rho^2}{|G|}, \tag{51}$$

regardless of whether or not there is a hidden shift. This is analogous to the fact that weak Fourier sampling is insufficient to distinguish between the trivial subgroup and the subgroups generated by full support involutions in the symmetric group [18].

Suppose we measure the representation name and observe a particular $\rho \in \hat{G}$. Then consider an arbitrary POVM $\mathcal{E} = \{a_1|\psi_1\rangle\langle\psi_1|, \ldots, a_r|\psi_r\rangle\langle\psi_r|\}$ acting on the subspace of dimension $2d_\rho$ corresponding to the observed representation.

If there is no hidden shift (that is, if the state is $\gamma_2^{(1)}$), then the post-measurement state is $I_{2d_\rho}/(2d_\rho)$, and the probability of obtaining the outcome j is

$$p_2(j) = \frac{a_j}{2d_\rho} \langle\psi_j|I_{2d_\rho}|\psi_j\rangle = \frac{a_j}{2d_\rho}. \tag{52}$$

We denote this probability distribution by P_2. On the other hand, if there is a hidden shift s, then the post-measurement state is $B^\rho(s)/(2d_\rho)$, and the probability of obtaining the outcome j is

$$p_1(j|s) := \frac{a_j}{2d_\rho} \langle\psi_j|B^\rho(s)|\psi_j\rangle; \tag{53}$$

we denote this distribution by $P_{1,s}$. We will also be interested in the distribution P_1 obtained by averaging over $s \in G$, i.e., with the probabilities

$$p_1(j) := \frac{1}{|G|} \sum_{s \in G} p_1(j|s). \tag{54}$$

12

Following [26, 27], the strategy for proving that single-register measurements are not sufficient is to show that with high probability (over the hidden shift s and the observed representation ρ), the statistics of the measurement results when there is a hidden shift s are close to those when there is no hidden shift. More precisely, we will prove

Theorem 8.
$$\Pr_{s \in G, \rho \in \hat{G}}(\|P_{1,s} - P_2\|_1 \geq e^{-\Theta(n)}) \leq e^{-\Theta(n)} \tag{55}$$

To prove this theorem, we first show that with high probability (over a uniformly random choice of $s \in G$ and the Plancherel distribution of irreducible representations $\rho \in \hat{G}$), the distribution $P_{1,s}$ is close to the distribution P_1. Then it suffices to show that P_1 and P_2 are typically close, which is straightforward (since in fact, they are typically identical).

Because P_1 is the average of $P_{1,s}$ over $s \in G$, we can show that the distributions are likely to be close by showing that the variance of $p_1(j|s)$ is small (so that we can apply the Chebyshev inequality). More precisely, we will use the following:

Lemma 9 (Upper bound on the sum of weighted variances). *Assume we have measured the irreducible representation $\rho \neq \hat{1}$, and we perform an arbitrary measurement $\mathcal{E} = \{a_j|\psi_j\rangle\langle\psi_j| : j \in J\}$. Then*
$$\sum_{j \in J} \frac{\sigma_j^2}{a_j} \leq \frac{1}{d_\rho^2} \tag{56}$$

where σ_j^2 is the variance of $p_1(j|s)$ when s is chosen uniformly from G.

Proof. For any fixed j the variance σ_j^2 is given by
$$\sigma_j^2 := \frac{1}{|G|} \sum_{s \in G} p_1(j|s)^2 - p_1(j)^2 . \tag{57}$$

Recall that we have $p_1(j) = a_j/(2d_\rho)$ for all j. This is because we have $B^\rho = I_{2d_\rho}$ for all $\rho \neq \hat{1}$ as shown in Lemma 6.

The second moment can be expressed in terms of the block $B^{\rho,\rho}$. We have

$$\frac{1}{|G|} \sum_{s \in G} p_1(j|s)^2 = \frac{a_j^2}{(2d_\rho)^2} \frac{1}{|G|} \sum_{s \in G} \left(\langle\psi_j|B^\rho(s)|\psi_j\rangle\right)^2 \tag{58}$$

$$= \frac{a_j^2}{(2d_\rho)^2} \frac{1}{|G|} \sum_{s \in G} \langle\psi_j|\langle\psi_j|B^\rho(s) \otimes B^\rho(s)|\psi_j\rangle|\psi_j\rangle \tag{59}$$

$$= \frac{a_j^2}{(2d_\rho)^2} \frac{1}{|G|} \sum_{s \in G} \langle\psi_j|\langle\psi_j|B^{\rho,\rho}(s)|\psi_j\rangle|\psi_j\rangle \tag{60}$$

$$= \frac{a_j^2}{(2d_\rho)^2} \langle\psi_j|\langle\psi_j|B^{\rho,\rho}|\psi_j\rangle|\psi_j\rangle . \tag{61}$$

Set $\Delta := |B^{\rho,\rho} - I|$. Then we have for the variance the upper bound

$$\sigma_j^2 \leq \frac{a_j^2}{(2d_\rho)^2} \langle\psi_j|\langle\psi_j|\Delta|\psi_j\rangle|\psi_j\rangle . \tag{62}$$

13

The operator Δ has the eigenvalue 1 occurring with multiplicity either 0 or 2 and the eigenvalue $1/d_\rho$ occurring with multiplicity $2d_\rho^2$. This follows from Lemma 7 where we have determined the spectrum of blocks of the form $B^{\rho,\rho}$. Denote the spectral decomposition of Δ by

$$\Delta = P + \frac{1}{d_\rho}Q \tag{63}$$

where P, Q are projectors. We bound the sum of the weighted variances by looking at P and Q/d_ρ separately. We have

$$\sum_{j \in J} a_j \langle \psi_j | \langle \psi_j | Q/d_\rho | \psi_j \rangle | \psi_j \rangle \leq \sum_{j \in J} \frac{a_j}{d_\rho} = 2 \,. \tag{64}$$

We also have

$$\sum_{j \in J} a_j \langle \psi_j | \langle \psi_j | P | \psi_j \rangle | \psi_j \rangle \leq \operatorname{rank} P \leq 2 \tag{65}$$

where the first inequality follows by Lemma 12 in [27]. Putting these two bounds together and multiplying by $1/(2d_\rho)^2$, we obtain the desired result. $\square$

Now we can use this result to show that $P_{1,s}$ and P_1 are probably close:

Lemma 10.

$$\Pr_{s \in G, \rho \in \hat{G}}(\|P_{1,s} - P_1\|_1 \geq e^{-\Theta(n)}) \leq e^{-\Theta(n)} \tag{66}$$

Proof. For any fixed representation $\rho \in \hat{G}$, according to Chebyshev's inequality,

$$\Pr_{s \in G}\left(\left|p_1(j|s) - p_1(j)\right| \geq a_j c\right) \leq \frac{\sigma_j^2}{a_j^2 c^2} \tag{67}$$

for any $c > 0$. Now define

$$J_{\text{bad}}^s := \left\{j \in J : \left|p_1(j|s) - p_1(j)\right| \geq a_j c\right\}, \tag{68}$$

and define $J_{\text{good}}^s := J - J_{\text{bad}}^s$. The total variation distance can be decomposed into contributions from good and bad j's. For the good j's, we have

$$\sum_{j \in J_{\text{good}}^s} \left|p_1(j|s) - p_1(j)\right| \leq \sum_{j \in J_{\text{good}}^s} a_j c \tag{69}$$

$$\leq 2d_\rho c \,. \tag{70}$$

Now for any $j \in J$ (and in particular, for $j \in J_{\text{bad}}^s$), we have

$$\left|p_1(j|s) - p_1(j)\right| = \frac{a_j}{2d_\rho}\left|\langle \psi_j | B^\rho(s) - B^\rho | \psi_j \rangle\right| \tag{71}$$

$$\leq \frac{a_j}{2d_\rho}\|B^\rho(s) - B^\rho\| \tag{72}$$

$$\leq \frac{a_j}{d_\rho} \,. \tag{73}$$

14

Thus it suffices to show that $\sum_{j \in J_{\mathrm{bad}}^s} a_j$ is small. The expectation of this quantity is

$$\mathbf{E}_{s \in G} \sum_{j \in J_{\mathrm{bad}}^s} a_j = \frac{1}{|G|} \sum_{s \in G} \sum_{j \in J} a_j \, \delta[j \in J_{\mathrm{bad}}^s] \tag{74}$$

$$= \sum_{j \in J} a_j \Pr_{s \in G}(j \in J_{\mathrm{bad}}^s) \tag{75}$$

$$\leq \sum_{j \in J} \frac{\sigma_j^2}{a_j c^2} \tag{76}$$

$$\leq \frac{1}{d_\rho^2 c^2} \tag{77}$$

where in the last line we have used Lemma 9 (assuming $\rho \neq \hat{1}$, which we will later ensure). Hence by Markov's inequality,

$$\Pr\left(\sum_{j \in J_{\mathrm{bad}}^s} a_j \geq c'\right) \leq \frac{1}{d_\rho^2 c^2 c'} \tag{78}$$

for any $c' > 0$. Conditioning on this event, we have

$$\|P_{1,s} - P_1\|_1 = \sum_{j \in J_{\mathrm{good}}^s} \left|p_1(j|s) - p_1(j)\right| + \sum_{j \in J_{\mathrm{bad}}^s} \left|p_1(j|s) - p_1(j)\right| \tag{79}$$

$$\leq 2 d_\rho c + \frac{c'}{d_\rho} \tag{80}$$

with probability at least

$$1 - \frac{1}{d_\rho^2 c^2 c'} \,. \tag{81}$$

Hence if we choose

$$c = \frac{e^{-\alpha n}}{d_\rho} \tag{82}$$

$$c' = e^{3\alpha n} \tag{83}$$

for some fixed $\alpha > 0$, we find

$$\|P_{1,s} - P_1\|_1 \leq 2 e^{-\alpha n} + \frac{e^{3\alpha n}}{d_\rho} \tag{84}$$

with probability at least

$$1 - e^{-\alpha n} \,. \tag{85}$$

For $P_{1,s}$ and P_1 to be close with high probability, it suffices that d_ρ is large with high probability, so that the second term of (84) is small. Thus we condition on the event that $d_\rho > n^{c'' n}$ for some constant c'', which occurs with probability at least $1 - n^{-\Omega(n)}$ [27, Lemma 6]. This completes the proof. $\qquad\square$

Finally, we must show that the probability distributions P_1 and P_2 are close in total variation distance:

Lemma 11. *For an arbitrary POVM acting on a single copy of the hidden shift state,*

$$\|P_1^\rho - P_2^\rho\| = 0 \tag{86}$$

for $\rho \neq \hat{1}$ and

$$\|P_1^{\hat{1}} - P_2^{\hat{1}}\| \leq \frac{1}{2} \tag{87}$$

for the trivial representation $\hat{1}$.

Proof. Let B be the block corresponding to the measured representation. Let $\Delta := |I - B|$. Then we have

$$\|P_1 - P_2\| = \frac{1}{2d_\rho} \sum_j a_j |\langle\psi_j|I_{d_\rho}|\psi_j\rangle - \langle\psi_j|B|\psi_j\rangle| \tag{88}$$

$$\leq \frac{1}{2d_\rho} \sum_j a_j \langle\psi_j|\Delta|\psi_j\rangle \tag{89}$$

$$= \frac{1}{2d_\rho} \sum_j \mathrm{tr}(a_j|\psi_j\rangle\langle\psi_j|\Delta) \tag{90}$$

$$= \frac{1}{2d_\rho} \mathrm{tr}(\Delta). \tag{91}$$

We have determined the spectrum of B in Lemma 6, from which the lemma follows. □

Putting these results together, we can now prove the main result:

Proof of Theorem 8. Since the trivial representation only appears with probability $1/n!$, we can simply condition on not obtaining the trivial representation, and the result follows from Lemmas 10 and 11. □

Acknowledgments

A Rank calculations

Although the measurement that projects on the support of $\gamma_1^{(k)}$ need not be optimal in general, it is nevertheless a natural measurement to consider—for example, an analogous measurement was used in [10] to show that $O(n \log n)$ hidden subgroup states are sufficient to solve a hidden subgroup problem relevant to graph isomorphism. Since we

are trying to distinguish $\gamma_1^{(k)}$ from the maximally mixed state, the success probability of the measurement that projects onto the support depends only on the rank of $\gamma_1^{(k)}$ (see (13)). Here we summarize some results on the rank for $k = 1$ and 2.

For the case $k = 1$, Lemma 6 immediately gives

$$\operatorname{rank} \gamma_1^{(1)} = 2|G| - 1 \,. \tag{92}$$

For the case $k = 2$, Lemma 7 gives the contribution to the rank from the cases where the same irreducible representation $\rho \neq \hat{1}$ occurs twice. It is straightforward to calculate the contribution from the other cases, giving the final result

$$\operatorname{rank} \gamma_1^{(2)} = 4|G|^2 - 6|G| + 3 + \sum_{\rho \in \hat{G}, \, d_\rho > 1} d_\rho^2 \tag{93}$$

$$= 4|G|^2 - 5|G| + 3 - |\{\rho \in \hat{G} : d_\rho = 1\}| \,. \tag{94}$$

In particular, for $G = S_n$, we have $|G| = n!$ and only two one-dimensional representations (the trivial and sign representations), so

$$\operatorname{rank} \gamma_1^{(2)} = 4(n!)^2 - 5n! + 1 \,. \tag{95}$$

Calculations of the rank for larger k would seem to require a better understanding of the structure of $\gamma_1^{(k)}$.

References

[1] N. Alon and J. H. Spencer, *The Probabilistic Method*, 2nd ed., Wiley Interscience, New York, 2000.

[2] L. Babai, W. M. Kantor, and E. M. Luks, *Computational complexity and the classification of finite simple groups*, Proceedings of the 24th Annual IEEE Symposium on Foundations of Computer Science, 1983, pp. 162–171.

[3] D. Bacon, A. M. Childs, and W. van Dam, *Optimal measurements for the dihedral hidden subgroup problem*, arXiv:quant-ph/0501044.

[4] _____, *From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups*, Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, 2005, arXiv:quant-ph/0504083.

[5] R. Beals, *Quantum computation of Fourier transforms over symmetric groups*, Proceedings of the 29th Annual ACM Symposium on Theory of Computing, ACM Press, New York, 1997, pp. 48–53.

[6] R. Bercov and L. Moser, *On Abelian permutation groups*, Canad. Math. Bull. **8** (1967), 627–630.

[7] R. Boneh and R. Lipton, *Quantum cryptanalysis of hidden linear functions*, Advances in Cryptology – Crypto'95, Lecture Notes in Computer Science, vol. 963, Springer-Verlag, Berlin, 1995, pp. 424–437.

[8] A. M. Childs and W. van Dam, *Quantum algorithm for a generalized hidden shift problem*, arXiv:quant-ph/0507190.

[9] W. van Dam, S. Hallgren, and L. Ip, *Quantum algorithms for some hidden shift problems*, Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, 2003, pp. 489–498, quant-ph/0211140.

[10] M. Ettinger and P. Høyer, *A quantum observable for the graph isomorphism problem*, arXiv:quant-ph/9901029.

[11] ———, *On quantum algorithms for noncommutative hidden subgroups*, Advances in Applied Mathematics **25** (2000), 239–251, arXiv:quant-ph/9807029.

[12] M. Ettinger, P. Høyer, and E. Knill, *The quantum query complexity of the hidden subgroup problem is polynomial*, Information Processing Letters **91** (2004), 43–48, arXiv:quant-ph/0401083.

[13] K. Friedl, G. Ivanyos, F. Magniez, M. Santha, and P. Sen, *Hidden translation and orbit coset in quantum computing*, Proceedings of the 35th Annual ACM Symposium on Theory of Computing, ACM Press, New York, 2003, pp. 1–9, arXiv:quant-ph/0211091.

[14] D. Gavinsky, *Quantum solution to the hidden subgroup problem for Poly-Near-Hamiltonian groups*, Quantum Information and Computation **4** (2004), 229–235.

[15] M. Grigni, L. Schulman, M. Vazirani, and U. Vazirani, *Quantum mechanical algorithms for the nonabelian hidden subgroup problem*, Combinatorica **24** (2004), 137–154.

[16] L. Hales and S. Hallgren, *An improved quantum Fourier transform algorithm and applications*, Proceedings of the 41st Annual Symposium on Foundations of Computer Science, IEEE, Los Alamitos, CA, 2000, pp. 515–525.

[17] S. Hallgren, M. Rötteler, and P. Sen, Personal communication.

[18] S. Hallgren, A. Russell, and A. Ta-Shma, *The hidden subgroup problem and quantum computation using group representations*, SIAM J. Comput. **32** (2003), 916–934.

[19] C. W. Helstrom, *Quantum Detection and Estimation Theory*, Academic Press, New York, 1976.

[20] P. Høyer, *Efficient quantum transforms*, arXiv:quant-ph/9702028.

[21] G. Ivanyos, F. Magniez, and M. Santha, *Efficient quantum algorithms for some instances of the non-abelian hidden subgroup problem*, International Journal of Foundations of Computer Science **14** (2003), 723–739, arXiv:quant-ph/0102014.

[22] J. Kempe and A. Shalev, *The hidden subgroup problem and permutation group theory*, Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2005, arXiv:quant-ph/0406046.

[23] A. Kitaev, *Quantum measurements and the abelian stabilizer problem*, arXiv:quant-ph/9511026.

[24] G. Kuperberg, *A subexponential-time quantum algorithm for the dihedral hidden subgroup problem*, arXiv:quant-ph/0302112.

[25] C. Moore, D. N. Rockmore, A. Russell, and L. J. Schulman, *The hidden subgroup problem in affine groups: Basis selection in Fourier sampling*, Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2004, pp. 1113–1122, arXiv:quant-ph/0211124, extended version at arXiv:quant-ph/0503095.

[26] C. Moore and A. Russell, *The symmetric group defies strong Fourier sampling: Part II*, arXiv:quant-ph/0501066.

[27] C. Moore, A. Russell, and L. J. Schulman, *The symmetric group defies strong Fourier sampling: Part I*, Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, 2005, arXiv:quant-ph/0501056.

[28] ———, Personal communication.

[29] E. Petrank and M. Roth, *Is code equivalence easy to decide?*, IEEE Trans. Inform. Theory **43** (1997), 1602–1604.

[30] O. Regev, *Quantum computation and lattice problems*, Proceedings of the 43rd Annual Symposium on Foundations of Computer Science, IEEE, Los Alamitos, CA, 2002, pp. 520–529, arXiv:cs.DS/0304005.

[31] _____, *A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space*, arXiv:quant-ph/0406151.

[32] P. W. Shor, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM Journal on Computing **26** (1997), 1484–1509.

[33] D. R. Simon, *On the power of quantum computation*, SIAM Journal on Computing **26** (1997), 1474–1483.

[34] N. J. A. Sloane, *The on-line encyclopedia of integer sequences*, 2005, `http://www.research.att.com/~njas/sequences`.

[35] A. M. Vershik and S. V. Kerov, *Asymptotic behavior of the maximum and generic dimensions of irreducible representations of the symmetric group*, Funct. Anal. Appl. **19** (1989), 21–31, English translation of Funk. Anal. i Prolizhen **19** (1985), no. 1, 25–26.

[36] P. Wocjan and M. Horodecki, *Characterization of combinatorially independent permutation separability criteria*, To appear in Open Systems and Information Dynamics. arXiv:quant-ph/0503129.